arXiv:1701.01547v1 [cs.RO] 6 Jan 2017

Stochastic Optimal Control for Modeling Reaching Movements in the Presence of Obstacles: Theory and Simulation

Arun Kumar Singh[1,2], *Member, IEEE*, Sigal Berman[1], *Member, IEEE*, and Ilana Nisky[2], *Member, IEEE*

***Abstract*— In many human-in-the-loop robotic applications such as robot-assisted surgery and remote teleoperation, predicting the intended motion of the human operator may be useful for successful implementation of shared control, guidance virtual fixtures, and predictive control. Developing computational models of human movements is a critical foundation for such motion prediction frameworks. With this motivation, we present a computational framework for modeling reaching movements in the presence of obstacles. We propose a stochastic optimal control framework that consists of probabilistic collision avoidance constraints and a cost function that trades-off between effort and end-state variance in the presence of a signal-dependent noise. First, we present a series of reformulations to convert the original non-linear and non-convex optimal control into a parametric quadratic programming problem. We show that the parameters can be tuned to model various collision avoidance strategies, thereby capturing the quintessential variability associated with human motion. Then, we present a simulation study that demonstrates the complex interaction between avoidance strategies, control cost, and the probability of collision avoidance. The proposed framework can benefit a variety of applications that require teleoperation in cluttered spaces, including robot-assisted surgery. In addition, it can also be viewed as a new optimizer which produces smooth and probabilistically-safe trajectories under signal dependent noise.**

I. Introduction

Motion prediction is important for finding the middle ground between pure teleoperation and autonomous control of robotic systems. It allows the robot to anticipate the future motions of the users and, consequently, their intention, and assist them in performing a given task. To improve the performance of motion prediction algorithms, it is beneficial to ground the prediction in experimentally-validated computational models of human movement [1]. Optimal control is used extensively in computational motor control, and provides a powerful framework for explaining a wide range of empirical phenomena associated with human motion [2], [3], [4]. In this view, it is hypothesized that human motion is driven by well-defined rewards or cost functions. The complimentary Inverse Optimal Control (IOC) framework attempts to identify the structure and parameters of these cost functions from a set of observed trajectories [7]. Thus, IOC allows for the transition from modeling of human motion to motion prediction in a particular task [8]. However, the accuracy of the model that is used for a particular problem is critical to the success of IOC-based approaches. While many studies considered optimal control for modeling reaching trajectories in free space [2], [4], [5], [6], there has been much less effort towards modeling reaching in the presence of obstacles using optimal control [22]. This in turn, hinders the development of efficient IOC based approaches for prediction.

In the current paper, we propose a stochastic optimal control framework for modeling human reaching trajectories in the presence of obstacles. This framework is designed to be incorporated in motion prediction for a variety of applications of teleoperation in cluttered spaces. Our proposed framework is built on experimental studies that suggest that reaching movements amongst obstacles are optimized considering the likelihood of collision [9], [10], [12], and that obstacle avoidance is sensitive to human perception of free space [10]. In line with these findings, the proposed optimal control model incorporates probabilistic collision avoidance constraints to ensure that the likelihood of collision is below a specified threshold. We also consider signal-dependent noise in human movement control [11], and the uncertainty in the perception of the size of the obstacle to model the error in estimation of free space.

We present a series of reformulations based on [14] to approximate a difficult non-linear and non-convex optimal control problem by a parametric quadratic optimization problem. The reformulations use substitution of chance constraints with a family of surrogate constraints. Satisfaction of each member of the family of surrogate constraints can be mapped to a lower bound probability with which the original chance constraints would be satisfied. Further, we show that the parameters of the reformulated quadratic optimization problem can be tuned to generate a diverse class of trajectories, and simulate the quintessential diversity of human motion. We present a simulation study highlighting the interaction of the parameters with the control cost and the probability of collision avoidance. We also discuss how the main insights derived from the simulation study agrees with the existing experimental findings on obstacle avoiding reaching movements.

The rest of the paper is organized as follows. Section II reviews the existing works which considers collision avoidance within the context of optimal control. Section III presents the optimal control problem followed by a series of reformulations to convert it into a tractable parametric quadratic optimization problem. Section IV presents a se-

The research was supported by the Helmsley Charitable Trust through the Agricultural, Biological and Cognitive Robotics Initiative of Ben-Gurion University of Negev, Israel, and by the Israeli Science Foundation (grant 823/15)

[1] Department of Industrial Engineering and Management, Ben-Gurion University of the Negev, Israel

[2] Department of Biomedical Engineering, Ben-Gurion University of the Negev, Israel

quential quadratic programming based algorithm for solving the reformulated optimization problem. Section V presents simulation results that demonstrate how the parameters of the reformulated problem result in a diverse set of trajectories and control costs. In section VI we discuss the results of our simulations in light of the existing experimental findings on reaching movements among obstacles and present future directions.

II. Related Work

Optimal control or optimization based obstacle avoidance has been extensively studied in the robotics literature, and to a smaller extent, in computational motor control.

A. Obstacle Avoidance in Robotics

Optimal control or optimization are used extensively to plan collision-avoiding trajectories that also optimize a specified cost function [15], [16]. In [16], optimal control is applied to stochastic systems with additive noise, and collision avoidance is ensured by introducing a penalty on trajectories that come close to the obstacles. An expectation over the cost is taken which suggests that the optimization is risk neutral; that is, it does not model the probability of collision avoidance. Trajectory optimizers like [17], [19] incorporates a penalty on the probability of collision avoidance.

Other studies put hard constraints on probability of collision avoidance [18], [20]. However, these studies assumed an additive noise model. We aim at planning trajectories for a human arm which is assumed to be modeled as a stochastic system with signal dependent noise [11]. An optimal control based framework presented in [21] presents collision avoidance under signal dependent noise, but for single integrator systems.

B. Obstacle Avoidance in Computational Motor Control

Reaching trajectories in the presence of obstacles were studied in computational motor control for understanding movement coordination. Experimental studies [9], [10], [12], [13] investigated the effects of obstacle position and size on obstacle avoidance. In particular, [13] observed that the obstacle avoidance strategy exhibited by human subject during reach to grasp movements, consisted of two basic but coupled components namely moving around the obstacle or slowing down near them.

An optimal control model for a single-obstacle avoidance was proposed in [22]. They solved the optimal control problem using simulated annealing. Simple obstacle configurations, predominantly with a single obstacle were considered. Our proposed approach differs from [22] in terms of the technical approach followed to solve the optimal control problem. In particular, we exploit some efficient structures in the problem. Moreover, we consider complex obstacle configurations to highlight the interaction between parameters, control cost and probability of collision avoidance.

III. Problem Formulation

A. Dynamics and Task Description

We consider the task of reaching movements in a 2D cluttered environment. We chose a simple linear model for the movement of the end point of the hand - a triple integrator – system. We denote the state of the hand at time instant t by $\mathbf{X}^t = (x^t, y^t, \dot{x}^t, \dot{y}^t, \ddot{x}^t, \ddot{y}^t)$, where the individual state variables are defined as the following normal distributions:

$$\begin{aligned} x^t &\approx N(\mu_x^t, \sigma_{x^t}^2), y^t \approx N(\mu_y^t, \sigma_{y^t}^2) \\ \dot{x}^t &\approx N(\mu_{\dot{x}^t}, \sigma_{\dot{x}^t}^2), \dot{y}^t \approx N(\mu_{\dot{y}^t}, \sigma_{\dot{y}^t}^2) \\ \ddot{x}^t &\approx N(\mu_{\ddot{x}^t}, \sigma_{\ddot{x}^t}^2), \ddot{y}^t \approx N(\mu_{\ddot{y}^t}, \sigma_{\ddot{y}^t}^2). \end{aligned} \tag{1}$$

The parameters of the distributions, i.e. their means and variances are obtained from the following discrete time dynamics with jerk $U = (u_x, u_y) = (\dddot{x}, \dddot{y})$ as the control input.

$$X^{t+1} = \mathbf{A}X^t + \mathbf{B}(U^t + \varepsilon_U^t), \tag{2}$$

where $\mathbf{A}$ and $\mathbf{B}$ represent state transition and control scaling matrices of dimensions conforming to that of the state, and

$$\varepsilon_U^t = \sum_{i=1}^{2} \phi_i \mathbf{M}_i U \tag{3}$$

$$\mathbf{M}_1 = \begin{bmatrix} c_x & 0 \\ 0 & 0 \end{bmatrix}, \mathbf{M}_2 = \begin{bmatrix} 0 & 0 \\ 0 & c_y \end{bmatrix}. \tag{4}$$

The term ε_U^t in (3) represents the time varying signal-dependent noise, and is formulated in terms of constant scaling matrices $\mathbf{M}_i$ and ϕ_i which are a set of zero-mean unit-variance normal random variables. This form of (3) ensures that indeed the standard deviation of the noise grows linearly with the magnitude of the control signal [4], and the constants c_x, c_y determine the magnitude of noise as a fraction of the control input.

B. Optimal Control

The discrete time optimal control can be represented by the following set of equations.

$$\begin{aligned} \min J_{opt} &= J_{U^t} + J_{X^t} \\ Pr(C_j^t(x^t, y^t, x_j, y_j, R_j) \leq 0) &\geq \eta, j = 1, 2..n, \end{aligned} \tag{5}$$

where

$$J_{U^t} = \|U\|^2, J_{X^t} = \sum_{t=t_0}^{t=t_f} E[L(X^t, U^t)], \tag{6}$$

$$L(X^t, U^t) = \sum_{i=1}^{6} w_i (X_i^t - X^{t_f})^2, \tag{7}$$

$$R_j \approx N(\mu_{R_j}, \sigma_{R_j}^2). \tag{8}$$

The objective function in (5) consists of a control effort term and a state-dependent term which penalizes the end point

variance of the trajectory. The term w_i determines the relative weighting between the components of the state-dependent cost term. The constraints $C_j(.) \leq 0$ in (5) represent the collision avoidance requirement in a deterministic setting. Thus, the set of inequalities in (5) signify constraints that the collision avoidance requirement is satisfied with a particular lower bound probability η. The terms x_j, y_j and R_j denote the position and size of the j^{th} obstacle. To model the uncertainty in the estimation of obstacle size, R_j is defined as normally-distributed random variable.

The optimization (5) is difficult to solve due to the constraints on probability of collision avoidance, also known as chance constraints, and are computationally intractable [23]. Hence, we reformulate these chance constraints into a tractable form and show that the reformulation naturally leads to an efficient optimization structure.

C. Reformulating Chance Constraints

We follow [14], and substitute of $Pr(C_i^t(.))$ with:

$$Pr(C_j^t(x^t, y^t, x_j, y_j, R_j) \leq 0) \geq \eta \tag{9}$$
$$\Rightarrow E[C_j^t(.)] + k\sqrt{(}Var[C_j^t(.)] \leq 0, \eta \geq \frac{k^2}{1+k^2}.$$

where $E[C_j^t(.)]$ and $Var[C_j^t(.)]$ represent the expectation and variance of the constraints $C_j^t(.)$ with respect to the random variables x^t, y^t. This suggests that satisfaction of the deterministic surrogate in 9 ensures satisfaction of the original probabilistic constraints with at least a probability $\frac{k^2}{1+k^2}$. In [14], it is shown that computing an analytical expression for $E[C_j^t(.)]$ and $Var[C_j^t(.)]$ in terms of random variable arguments x^t, y^t, R_j etc. is simpler compared to computing that for $Pr(C_j^t(.))$.

We can further simplify (9) by approximating obstacle regions in 2D as circles. This simplifies the collision avoidance inequality $C_j^t(.)$:

$$C_j^t : -(x^t - x_j)^2 - (y^t - y_j)^2 + R_j^2 \leq 0. \tag{10}$$

Because (10) is purely concave in terms of hand position variables x^t and y^t, an affine upper bound can by obtained by linearizing C_i^t around an initial trajectory guess (x_*^t, y_*^t) [24]:

$$C_j^t \approx^* C_j^t + \bigtriangledown_{x^t} C_j^t (x^t - x*^t) + \bigtriangledown_{y^t} C_j^t (y^t - y*^t) \leq 0, \tag{11}$$

where ${}^*C_j^t$ is obtained by evaluating (10) at (x_*^t, y_*^t). Similarly, $\bigtriangledown_{x^t} C_j^t$ and $\bigtriangledown_{y^t} C_j^t$ represent the partial derivative of $C_j^t(.)$ with respect to x^t and y^t, evaluated at (x_*^t, y_*^t). The affine approximation (11) can be further improved by updating (x_*^t, y_*^t), during the course of the optimization. This sequential linearization of concave constraints forms the basis of the *convex concave procedure* [24].

In light of (11), $E[C_j^t(.)]$ and $Var[C_j^t(.)]$ take the form

$$E[C_j^t(.)] = \sigma_{R_j}^2 \tag{12}$$
$$+h_1(\mu_{x^t}, x_*^t, \mu_{y^t}, , y_*^t, \sigma_{x^t}^2, \sigma_{y_t}^2 \mu_{x_j}, \mu_{y_j}, \mu_{R_j})$$

$$Var[C_j^t(.)] = C_{R_j} \sigma_{R_j}^2 + 2\sigma_{R_j}^4 \tag{13}$$
$$+h_2(\mu_{x^t}, x_*^t, \mu_{y^t}, y_*^t, \sigma_{x^t}^2, \sigma_{y^t}^2, \mu_{x_j}, \mu_{y_j}, \mu_{R_j}),$$

where the terms (μ_{x^t}, μ_{y^t}) and $(\sigma_{x^t}^2, \sigma_{y^t}^2)$ represent the mean and variance of the hand position (x^t, y^t). The term C_{R_j} and functions $h_1(.)$ and $h_2(.)$ are given below. It can be noted that $h_2(.)$ can be represented as sum of squares and thus, is non-negative.

$$C_{R_i} = 4\mu_{R_i}^2 \tag{14}$$

$$h_1 = \mu_{R_i}^2 + 2\mu_{x^t}\mu_{x_i} - \mu_{x_i}^2 + 2\mu_{y^t}\mu_{y_i} - \mu_{y_i}^2 - 2\mu_{x^t}x_*^t \tag{15}$$
$$-2\mu_{y^t}y_*^t + (x_*^t)^2 + (y_*^t)^2$$

$$h_2 = 2(2\mu_{x_i}^2\sigma_{x^t}^2 + 2\mu_{y_i}^2\sigma_{y^t}^2 - 4\mu_{x_i}\sigma_{x^t}^2(x_*^t)^2 - 4\mu_{y_i}\sigma_{y^t}^2(y_*^t)^2 \tag{16}$$
$$+2\sigma_{x^t}^2(x_*^t)^2 + 2\sigma_{y^t}^2(y_*^t)^2)$$

D. Reformulated Optimal Control Problem

To arrive at the final reformulated version of (5), we make the following sequence of observations. The second term of the surrogate constraints proposed in (9) is non-negative. Thus, for a given k, the surrogate constraints (9) are satisfied when the first term, $E[C_j^t(.)]$ is sufficiently negative and the second term, $\sqrt{(}Var[C_j^t(.)]$ is sufficiently small in magnitude. Due to (13) and (16) we note that $\sqrt{(}Var[C_j^t(.)]$ is a non-decreasing function of the positional variance at each point of the trajectory $(\sigma_{x^t}^2, \sigma_{y^t}^2)$. Thus, making $\sqrt{(}Var[C_j^t(.)]$ small is equivalent to minimizing the positional variance at each point of the trajectory. In light of all these arguments, FOC (5) can be replaced with the following simpler problem.

$$J_{aug} = \|U\|^2 + \sum_{t=t_0}^{t=t_f} E[L(X^t, U^t)] + \lambda \sum_{t=t_0}^{t_f} (\sigma_{x^t}^2 + \sigma_{y^t}^2) \tag{17}$$
$$E[C_j^t(.)] + \tau \leq 0$$

The original trajectory optimization (5) has been converted to the new formulation (17) by substituting the parameter η which represented probability of avoidance in (5) with two new sets of variables τ and λ. The positive constant τ can be manipulated to make $E[C_j^t(.)]$ as negative as required and consequently control the clearance from a given set of obstacles. Similarly, λ is a positive constant which can be manipulated to minimize the positional variance at each point along the trajectory. Hence, we can manipulate τ and λ to achieve a particular probability of avoidance η. Moreover, each η can be mapped to various choices of τ and λ leading to a diverse set of collision avoidance behaviors. Within this diverse set, τ determines the geometry of the path, and λ determines the velocity profile along the path.

The reformulated FOC (17) is very different from those typically used in the context of human motion modeling. A central hypothesis in current frameworks is that relative weighting of each term in the cost function can be tuned to produce a diverse set of trajectories. The FOC (17) takes on

a different approach – its parameters appear not only in the cost function but also in the constraints.

The reformulated FOC (17) can be solved in one shot if the right set of τ and λ are given. For the cases where such set is not available, we present a framework for mapping a probability of collision avoidance η to τ and λ and solving (17) in the process. This is discussed next.

IV. Efficiently Solving the Proposed FOC

Algorithm 1 summarizes a sequential quadratic programming (SQP) routine for solving FOC (17). The optimization starts with an initial guess trajectory (x_*^t, y_*^t) and initialization of an index counter i and two non-negative variables τ and λ. The outermost loop checks whether the constraints are satisfied and reduction in the cost function between two consecutive iterations is within a specified threshold, ξ. If either of these checks are violated, then the algorithm proceeds to the inner loop where we check whether the surrogate constraints (9) are satisfied. If not, then we increment the value of the τ by δ and λ by a factor of Δ. Thereafter (17) is solved with the current values of τ and λ and the solution obtained is used to update the initial guess trajectory, which in turn is used to obtain a better estimate of $C_j^i(.)$ through (11) for the next iteration.

Algorithm 1 has two important features. Firstly, $E[C_j^t(.)]$ is affine and J_{aug} is convex quadratic in terms of control variables. Thus, solving (17) for a given τ and λ amounts to solving a quadratic programming (QP) problem. This is turn can be accomplished efficiently through open source solvers like CVX [25]. Secondly, algorithm 1 is different from the standard SQP routines used to solve general non-convex problems in the sense that it does not require a trust region update. This, in turn, is because the affine approximation of C_j^t in (11) acts as a global upper bound for the original collision constraints (10)

Each η can be mapped to numerous combinations of τ and λ. This redundancy is captured in algorithm 1 by manipulating the update rates of τ and λ. We discuss this in more detail in Section V with the help of specific examples.

Algorithm 1 Sequential Quadratic Programming for solving FOC

Initialization: Initial guess for optimal trajectory x_*^t, y_*^t.
$i = 0$,$\tau = 0$, $\lambda = 1$

while **do**$|J_{opt}^{i+1} - J_{opt}^{i}| < \xi$ and $E[C_j^t(.)] + k\sqrt{(}Var[C_j^t(.)] \leq 0$
  if $E[C_j^t(.)] + k\sqrt{(}Var[C_j^t(.)] > 0$ **then**
$\tau \leftarrow \tau + \delta$
$\Delta \leftarrow \Delta\lambda$

  end if
$U \leftarrow \arg\min J_{aug}$
  $E[C_j^t(.)] + \tau \leq 0$
  Update x_*^t, y_*^t through U
$i \leftarrow i + 1$
end while

V. Simulation Results

A. Collision Avoidance Strategies

To ensure collision avoidance, a human can choose to maintain high clearance from the obstacles resulting in a large deviation from straight line paths. Alternatively, it can choose to reduce the deviation but compensate for it by moving with high precision near the obstacles (reduce positional variance), which in light of the the signal dependent noise (3), requires moving with low velocities near the obstacles. For the ease of exposition, from hereon, we will refer to the slowing down strategy as "Low Velocity" or **LV** and strategy of maintaining large clearance from the obstacles as "high clearance" or **HC**.

Both these strategies can be modeled through (17) by using appropriate values for parameters τ and λ. For example, Fig. 1 shows two solution trajectories of (17) between the same start and goal configurations. The probability of avoidance, η for both trajectories is 0.94. However, both trajectories achieve this probability of collision avoidance through different combinations of τ and λ. The trajectory resulting from strategy **LV** was obtained with $\tau = 0.0009, \lambda = 2.28 * 10^6$, while that resulting from strategy **HC** was obtained with $\tau = 0.0012, \lambda = 0.9 * 10^6$. These values of τ and λ were obtained using different update rates of of τ and λ in algorithm 1. For simulating strategy **LV** we used $\delta = 0.00005$, $\Delta = 10$ in the update rule of τ and λ, and for simulating strategy **HC** we used $\delta = 0.0001$, $\Delta = 10$. Since, τ controls the clearance from the obstacles, setting higher update rates for τ resulted in trajectories belonging to strategy **HC**. On the other hand, a lower update rate for τ puts a higher emphasis on λ and consequently manipulation of positional variance through velocity control for collision avoidance, thus, resulting in trajectories belonging to strategy **LV**.

The velocity profiles shown in Fig. 1(b) demonstrate that a higher λ forces the velocity magnitude along the trajectory closer to the obstacle (strategy **LV**) to be small during the initial stages, i.e, while the trajectory is near the obstacles. Consequently, the positional variance is reduced and desired probability of collision avoidance is maintained. In contrast, the trajectory with higher clearance from the obstacle (strategy **HC**) has the liberty to move with faster velocity and let the variance of the movement grow. The velocity magnitude along trajectory resulting from strategy **LV** increases eventually, but this happens towards the end of the movement, after crossing the obstacles.

B. Mapping Avoidance Strategies to Control Cost

If we would derive a variant of the optimization (17) for a system with an additive constant-variance noise, the probability of collision avoidance, η would solely depend on the clearance from the obstacles. Thus, increase in η would directly lead to an increase in arc lengths, and consequently, control costs. However, to develop a framework that is suitable for modeling human arm movements, we

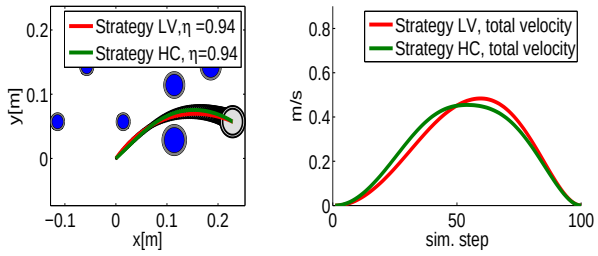

(a) (b)

Fig. 1. Demonstration of the effect of the choice of τ and λ on the collision avoidance strategies. Two sets of trajectories between same start and goal locations and having same probability of avoidance, η were computed. However, to generate these two trajectories we used a different set of τ and λ to achieve the specified probability of avoidance. The trajectories shown in green were computed using $\tau = 0.0012, \lambda = 0.9 * 10^6$, while trajectories shown in red were computed using $\tau = 0.0009, \lambda = 2.28 * 10^6$.

incorporated signal dependent noise [11]. In the presence of signal-dependent noise, control cost of trajectories depends on the probability of avoidance η, and more importantly, on the combination of τ and λ that is used in the optimization (17) to achieve this η. In other words, the control cost depends on the strategy that is used to achieve a particular probability of collision avoidance.

In Fig. 2(a)-(2(d)) we present simulated trajectories that correspond to both strategy **LV** and **HC** for probabilities of collision avoidance $\eta = 0.86$ and $\eta = 0.95$. The paths that resulted from strategy **HC** indeed has higher clearance from the obstacles. In contrast, the paths that resulted from strategy **LV** have lower clearance and thus, heavily rely on modifying the velocity magnitudes and consequently positional variance for collision avoidance. Consequently, paths resulting from strategy **HC** have higher arc lengths as compared to paths resulting from strategy **LV**. In Fig. (4(e)) the ratio of control costs for trajectories resulting from both the strategies is presented as a function of η. For low η, paths resulting from strategy **LV** which have lower arc lengths are less costly. But, as η increases, the higher arc length paths resulting from strategy **HC** become less costly.

The observations discussed above are apparent from the structure of the optimization (17). Increasing either τ, λ, or both, leads to an increase in the control cost. At low values of η, there is very little restriction on the growth of positional variance and thus the control cost is dictated by τ which controls the arc length. But as η increases, the effect of λ becomes prominent. This is consistent with the significant reduction in positional variance that is depicted in Fig. 2(c) and the corresponding skewed velocity profile shown in Fig. 2(d). Since trajectories resulting from strategy **LV** has lesser clearance from the obstacles, they require a higher value of λ to achieve the same η (similar to the result shown in previous section). Thus, at higher probabilities trajectories resulting from strategy **LV** become more costly in spite of having lower arc lengths.

The results presented above, were obtained with $c_x = c_y = 0.15$ in (3). That is, the noise was 15% of the control

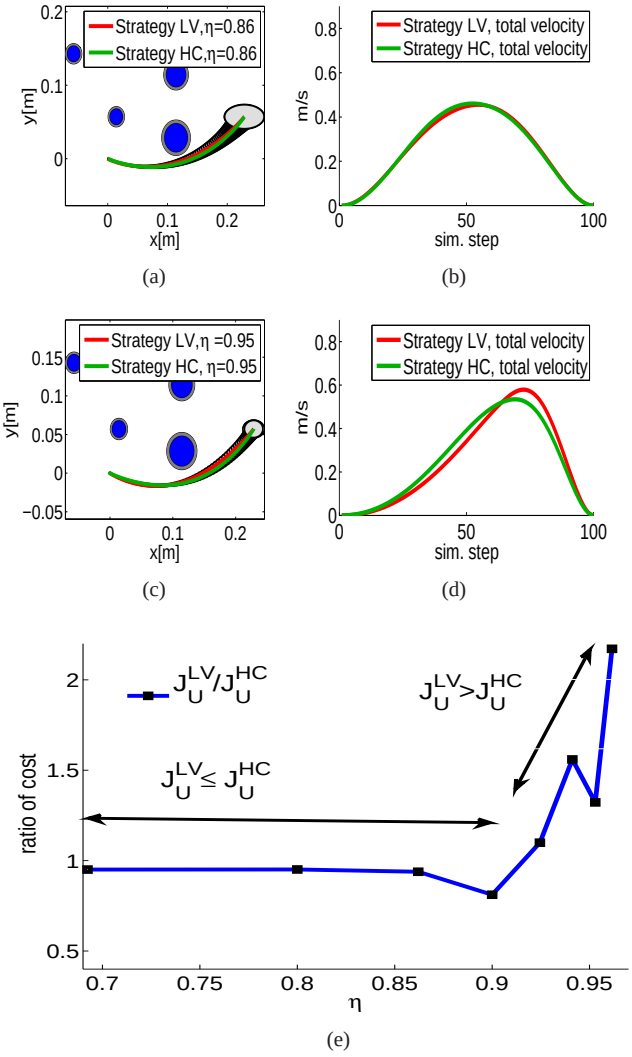

(a) (b) (c) (d) (e)

Fig. 2. Control costs vary with probability of avoidance. (a)- (d) Movements with different strategies between the same start and goal locations, the same obstacle configurations, and with noise level $c_x, c_y = 0.15$. (a), (c) present the paths with standard deviation ellipses of the two strategies. The obstacles are represented as blue filled circles and grey shaded region around them represent uncertainty about the size of the obstacle. (b), (d) present the velocity profiles. (e) the ratio of the control costs between the two strategies, $\frac{J_U^{LV}}{J_U^{HC}}$, as a function of η.

input. Next, we examined how the cost shown in Fig. 2(e) changes with a reduction in noise. Fig. 3 depicts the ratio of control costs for trajectories resulting from strategy **LV** and **HC** for $c_x = c_y = 0.05$. With lower noise, strategy **LV** becomes less costly even for higher probabilities. This result agrees with the common intuition. With a lesser noise there is no need to ensure high clearance from the obstacles, thereby making strategy **HC** redundant. In fact, for a zero noise system, the trajectory with least control cost would just graze the obstacle.

We would like to highlight that Fig. 2(e) and Fig. 3 are intended to demonstrate the general trend in ratio of control costs. An in depth analysis of the exact values and their dependence on the initial conditions of the optimization are

beyond the scope of this current study.

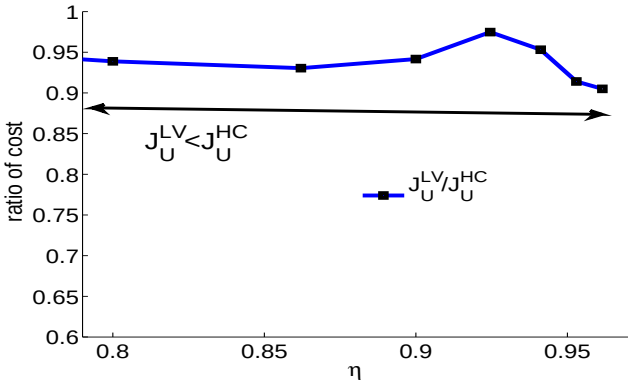

Fig. 3. The ratio of the control costs between the two strategies, $\frac{J_U^{LV}}{J_U^{HC}}$, as a function of η for noise level of $c_x, c_y = 0.05$.

C. Modeling Choice of Homotopies

In this section, we discuss how choice of strategy of collision avoidance or in other words, choice of τ and λ for a given η affects control cost of trajectories in different homotopies.

*1) Strategy **LV**:* In Fig. 4(a) and (4(c)) solution trajectories of (17) having same start and goal positions, but belonging to different homotopies and having different probability of avoidance, η, are depicted. The trajectories in both the homotopies were generated by choosing such values for τ and λ that ensure collision avoidance by slowing down near the obstacles and reducing positional variance (strategy **LV**) rather than taking a large deviation from them. Thus, as η increases from 0.9 (figure 4(a)) to 0.965 (figure 4(c)), we observe only a small change in arc length, but a significant change in the positional variance along the trajectories. Moreover, since trajectories of homotopy 2 move through a more cluttered environment, the reduction of positional variance along it is higher than that along trajectories of homotopy 1.

It is possible to relate the change in positional variance as η increases to the change in the control costs through the velocity profiles. Firstly, in contrast to Fig. 4(b), velocity profiles shown in Fig. 4(d) are skewed; i.e, they have low magnitudes during the initial phases and a peak which is shifted towards the right. This is to ensure that velocity magnitudes (and thus positional variance) are low near the obstacles and reach peak only after crossing the obstacles. Since trajectories in homotopy 2 require a larger reduction in positional variance, the skewness observed in their velocity profile is also higher. Finally, the skewness in velocity profiles is accompanied with higher peak velocities. This is because of the fixed final time paradigm of the optimization, (17). Since, magnitudes are low during initial phases of the trajectories, it needs to be compensated by moving faster in obstacle free space to ensure that the goal position is reached in specified time. Now, it is easy to deduce that a skewed velocity profile with higher peaks would mean higher accelerations and jerks and thus, consequently higher control costs.

To summarize, for collision avoidance strategy **LV**, maintaining high η requires larger reduction in positional variance leading to larger skewness in velocity profiles and consequently higher control costs. However, since trajectories in homotopy 2 require a larger reduction in positional variance, the control costs along it would increase at a higher rate than that along trajectories in homotopy 1. We demonstrate this last observation in Fig. 4(e) which shows the ratio of control costs along homotopy 1 and homotopy 2 for the various values of η. For lower values (till $\eta = 0.9$), the cost along homotopy 1 and homotopy 2 are similar owing to their similar velocity profiles. However, for higher η, cost along homotopy 1 is significantly lower than that along homotopy 2.

*2) Strategy **HC**:* Here we re-analyze the cost along homotopies for the same configuration as shown in Fig. 4, but with respect to strategy **HC** where there is a bigger reliance on clearance from the obstacles for collision avoidance.

The trajectories along both the homotopies are shown in Fig. 5(a) and 5(c). Comparing these trajectories with Fig. 4(a) and 4(c), demonstrates that there is a significant increase in clearance from the obstacles with increase in η. Thus, a lesser restriction is required on the growth of positional variance and consequently, the velocity profiles along trajectories in both the homotopy become very similar even at higher η (figure 5(d)). This is very different from the comparisons shown in Fig. 4(d).

The similarity in velocity trajectories in turn results in similar control costs along both the homotopies (figure 5(e)). In particular, the control cost along homotopies 1 and 2 are similar for a larger range of η. The lowest ratio of cost is 0.67 in figure 5(e) for $\eta = 0.9615$. In comparison, the ratio was 0.33 in figure 4(e) for the same η.

VI. Discussion and Future Work

In this paper, we presented a stochastic optimal control problem with signal dependent noise and probabilistic collision avoidance constraints as a model of human reaching among obstacles. We then reformulated it into a parameter optimization problem. The parameters τ and λ which appeared in the reformulated optimization problem, (17) served as a mapping between the probability of collision avoidance, η, and possible collision avoidance strategies.

The parameter τ models the clearance from the obstacles, and the parameter λ models the effect of slowing down near the obstacles. In our simulations, we demonstrate that effect of these two parameters on movement paths and velocity profiles is in agreement with the experimental findings reported in [13] for reach to grasp movements around obstacles. Specifically, they discussed two basic but coupled strategies of collision avoidance which consists of moving around the obstacle and slowing down.

We showed how each avoidance strategy results in a unique variation of control costs with respect to η both within and across homotopies. These variation in control

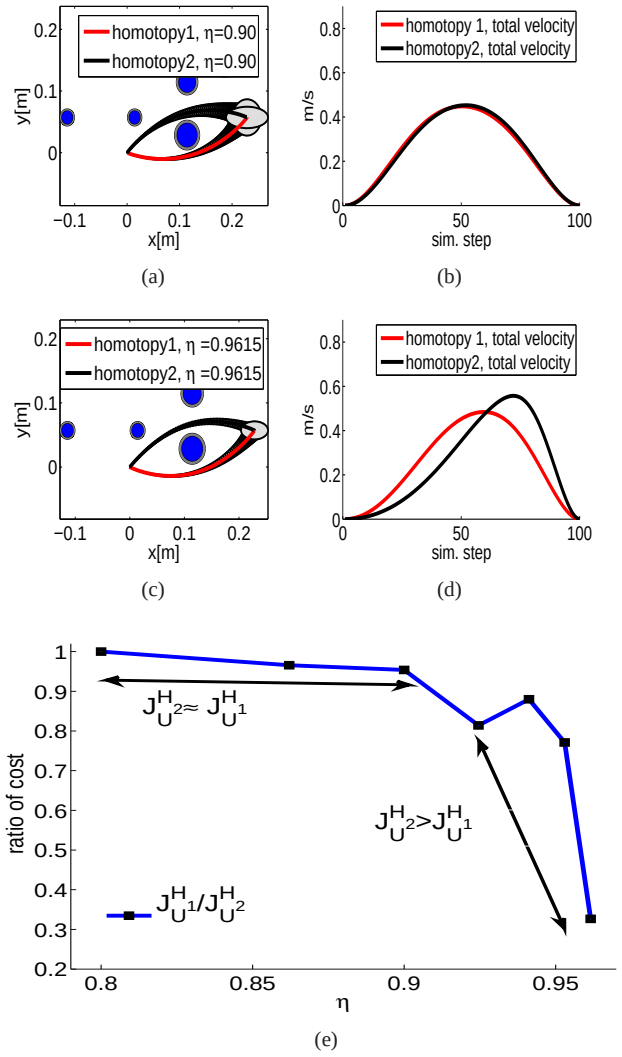

Fig. 4. Movements between the same start and goal locations and obstacle configurations but with different probability of avoidance. (a), (c) present the paths with standard deviation ellipses of the two homotopies. The obstacles are represented as blue filled circles and grey shaded region around them represent uncertainty about the size of the obstacle. (b), (d) present the velocity profiles. (e) the ratio of the control costs between the two homotopies, $\frac{J_U^{H_1}}{J_U^{H_2}}$, as a function of η. For the chosen avoidance strategy **LV**, the control cost along the homotopies is similar for low η. For higher η, the control cost along homotopy 1 is significantly less.

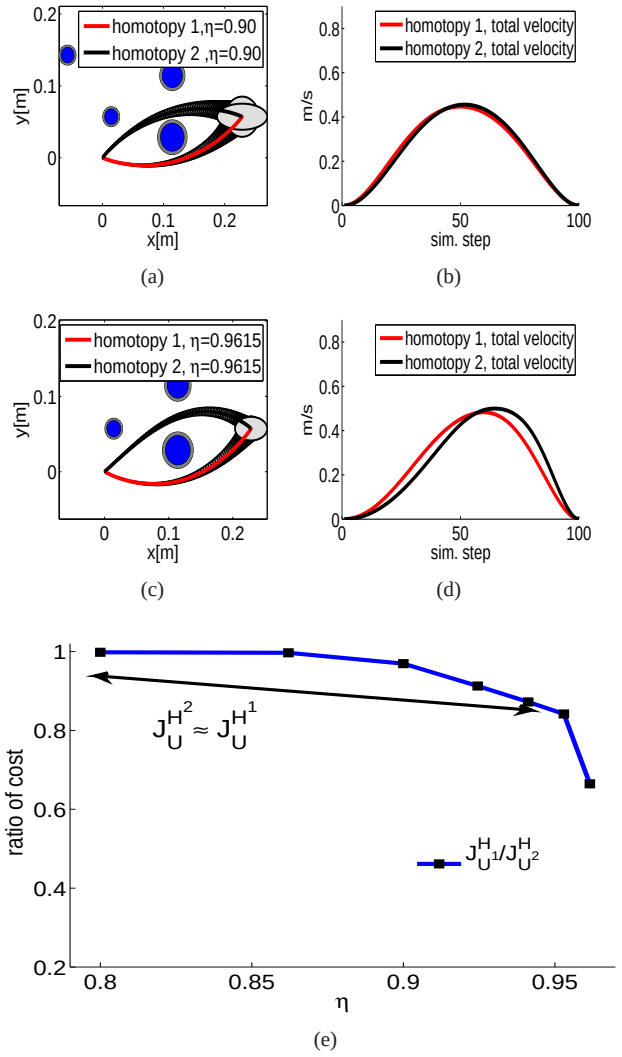

Fig. 5. Movements between the same start and goal locations and obstacle configurations but with different probability of avoidance, η . The results are similar to that shown in Fig. 4, but trajectories are now computed with respect to strategy **HC**, which gives higher emphasis on clearance from obstacles for obstacle avoidance. (e): Ratio of control cost along the two homotopies, $\frac{J_U^{H_1}}{J_U^{H_2}}$ with respect to strategy **HC**.

costs can be used as a basis for predicting user behavior between a given start and goal position and for a given obstacle environment. For example, in Fig. 2(a)-2(c), 2(e), we showed that a risk-seeking behavior (low η) is more likely to use strategy **LV** for collision avoidance as it requires less control effort. In contrast, a risk-averse behavior (high η) would likely choose strategy **HC**.

We also showed how control cost along different homotopies is dictated by the choice of avoidance strategy. This variation in control cost can be used to predict the homotopy selection by the human. In particular, if two competing homotopies have similar control costs, then the human would have equal affinity towards either of it, thus leading to a random behavior. However, as the ratio of control costs departs from unity, the possibility of selection would incline towards the lesser control costs, thus leading to more well defined pattern.

Our proposed framework has the following limitations. Firstly, we used a very simple dynamic model, and thus, we necessarily do not capture every aspect of the motion of the human arm. A second order linear mechanical system or a non-linear model of a serial link robotic manipulator are a better alternative. The second order mechanical system can be easily incorporated because as long as the system is linear, the structure of the optimization (17) would not change. In contrast, incorporating even a simple planar two-link manipulator model is more challenging, and may require methods similar to that proposed in [26]. Secondly, the fixed

final time paradigm of optimization (17) is not equipped to capture the effect of increase in traversal time of reaching motions due to presence of obstacles. A possible solution to this could be developed using the time scaling concepts [14].

From our simulation study, we conclude that if the parameters τ and λ are known, the possible choice of homotopy as well as choice of trajectory within that homotopy can be predicted. Thus, currently our efforts are focused towards developing an inverse optimization framework which can automatically recover these parameters from example trajectories demonstrated by the human. Furthermore, our simulation yielded testable predictions about homotopy and strategy choices that we plan to confirm in human experiments. These two advancements taken together will eventually lay the ground towards using our framework for motion prediction in realistic scenarios.

REFERENCES

[1] Wolpert, Daniel M., and Zoubin Ghahramani. "Computational principles of movement neuroscience." *Nature Neuroscience* 3 (2000): 1212-1217.

[2] Flash, Tamar, and Neville Hogan. "The coordination of arm movements: an experimentally confirmed mathematical model." *The journal of Neuroscience* 5, no. 7 (1985): 1688-1703.

[3] Flash, Tamar, Yaron Meirovitch, and Avi Barliya. "Models of human movement: Trajectory planning and inverse kinematics studies." *Robotics and Autonomous Systems* 61, no. 4 (2013): 330-339.

[4] Todorov, Emanuel. "Optimality principles in sensorimotor control." *Nature Neuroscience* 7, no. 9 (2004): 907-915.

[5] Scott, Stephen H. "Optimal feedback control and the neural basis of volitional motor control." *Nature Reviews Neuroscience* 5, no. 7 (2004): 532-546.

[6] Diedrichsen, Jrn, Reza Shadmehr, and Richard B. Ivry. "The coordination of movement: optimal feedback control and beyond." *Trends in cognitive sciences* 14, no. 1 (2010): 31-39.

[7] Li, Weiwei, Emanuel Todorov, and Dan Liu. "Inverse optimality design for biological movement systems." *IFAC Proceedings* Volumes 44, no. 1 (2011): 9662-9667.

[8] Mainprice, Jim, Rafi Hayne, and Dmitry Berenson. "Predicting human reaching motion in collaborative tasks using inverse optimal control and iterative re-planning." In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 885-892. IEEE, 2015.

[9] Chapman, Craig S., and Melvyn A. Goodale. "Missing in action: the effect of obstacle position and size on avoidance while reaching." *Experimental Brain Research* 191, no. 1 (2008): 83-97.

[10] Chapman, Craig S., and Melvyn A. Goodale. "Seeing all the obstacles in your way: the effect of visual feedback and visual feedback schedule on obstacle avoidance while reaching." *Experimental Brain Research* 202, no. 2 (2010): 363-375.

[11] Harris, Christopher M., and Daniel M. Wolpert. "Signal-dependent noise determines motor planning." *Nature* 394, no. 6695 (1998): 780-784.

[12] Mon-Williams, M., and R. D. McIntosh. "A test between two hypotheses and a possible third way for the control of prehension." *Experimental Brain Research* 134, no. 2 (2000): 268-273.

[13] Tresilian, James R. "Attention in action or obstruction of movement? A kinematic analysis of avoidance behavior in prehension." *Experimental Brain Research* 120, no. 3 (1998): 352-368.

[14] Gopalakrishnan, Bharath, Arun Kumar Singh, and K. Madhava Krishna. "Closed form characterization of collision free velocities and confidence bounds for non-holonomic robots in uncertain dynamic environments." In Proc. of *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pp. 4961-4968. IEEE, 2015.

[15] Zucker, Matt, Nathan Ratliff, Anca D. Dragan, Mihail Pivtoraiko, Matthew Klingensmith, Christopher M. Dellin, J. Andrew Bagnell, and Siddhartha S. Srinivasa. "CHOMP: Covariant Hamiltonian optimization for motion planning." *The International Journal of Robotics Research* 32, no. 9-10 (2013): 1164-1193.

[16] Kalakrishnan, Mrinal, Sachin Chitta, Evangelos Theodorou, Peter Pastor, and Stefan Schaal. "STOMP: Stochastic trajectory optimization for motion planning." In Proc. of *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 4569-4574. IEEE, 2011.

[17] Schulman, John, Yan Duan, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow, Jia Pan, Sachin Patil, Ken Goldberg, and Pieter Abbeel. "Motion planning with sequential convex optimization and convex collision checking." *The International Journal of Robotics Research* 33, no. 9 (2014): 1251-1270.

[18] Du Toit, Noel E., and Joel W. Burdick. "Robot motion planning in dynamic, uncertain environments." *IEEE Transactions on Robotics* 28, no. 1 (2012): 101-115.

[19] Mller, Jrg, and Gaurav S. Sukhatme. "Risk-aware trajectory generation with application to safe quadrotor landing." In Proc. of *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3642-3648. IEEE, 2014.

[20] Blackmore, Lars, Masahiro Ono, and Brian C. Williams. "Chance-constrained optimal path planning with obstacles." *IEEE Transactions on Robotics* 27, no. 6 (2011): 1080-1094.

[21] Van Den Berg, Jur, Sachin Patil, and Ron Alterovitz. "Motion planning under uncertainty using iterative local optimization in belief space." *The International Journal of Robotics Research* 31, no. 11 (2012): 1263-1278.

[22] Hamilton, Antonia F. de C., and Daniel M. Wolpert. "Controlling the statistics of action: obstacle avoidance." *Journal of neurophysiology* 87, no. 5 (2002): 2434-2440.

[23] Nemirovski, Arkadi. "On safe tractable approximations of chance constraints." *European Journal of Operational Research* 219, no. 3 (2012): 707-718.

[24] Boyd, Stephen. "Sequential convex programming." *Lecture Notes, Stanford University* (2008).

[25] Grant, Michael, and Stephen Boyd. "CVX: Matlab software for disciplined convex programming, version 1.21 (2011)." *Available: cvxr. com/cvx (2010)*.

[26] Todorov, Emanuel, and Weiwei Li. "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems." In Proc. of *2005, American Control Conference, 2005*, pp. 300-306. IEEE, 2005.